

Claims

This listing of claims will replace all prior versions and listings of claims in the application.

Listing of Claims:

1. (Currently Amended): One or more processor-accessible storage media comprising processor-executable instructions that, when executed, direct a first device to perform actions comprising:

accepting a connection from a connecting device at a forwarder;
receiving data at the forwarder from the connecting device as a result of accepting the connection;
forwarding the data from the forwarder to a classifier;
determining, by the classifier, a second device for receiving the connection;
aggregating a connection state for the connection at the classifier by aggregating a protocol state of a first protocol stack and the data to constitute a binary blob;
sending the connection state from the classifier to the second device for injection into a second protocol stack at the second device by sending the binary blob including the protocol state and the data to the second device, whereby the connection is transferred to the second device, wherein the sending the connection state further comprises transmitting the binary blob from the classifier to the second device asynchronously via the forwarder in a reliable manner such that the binary blob is

received intact at the second device even if one or more packets that comprise the binary blob are lost or corrupted;

in conjunction with sending the connection state, adding an entry to a mapping table maintained by the forwarder that indicates the second device as a destination for packets for the connection;

sending a mapping for a flow identifier to the second device based upon the entry in the mapping table;

receiving subsequent communications from the connecting device by the forwarder; and

encapsulating the subsequent communications by the forwarder according to the entry in the mapping table of the forwarder by inserting the flow identifier into the encapsulated communications; and receiving the encapsulated communications at the second device from the forwarder, wherein the flow identifier serves to identify a flow of encapsulated communications received at the second device from the forwarder as being associated with the connection to the connecting device.

2. (Previously Presented): The one or more processor-accessible storage media as recited in claim 1, further comprising, prior to the aggregating,

determining, by the classifier, the second device to receive migration of the connection state from among a plurality of second devices; and

passing a migrate connection function call to a topmost layer of the first protocol stack to initiate the aggregating of the connection state for migrating the connection state to the determined second device.

3. (Previously Presented): The one or more processor-accessible storage media as recited in claim 1, wherein the action of sending the connection state comprises an action of:

sending the binary blob including the protocol state and the data asynchronously to a connection migrator component at the second device, wherein the connection migrator component is configured to receive the binary blob as a bundle, reassemble the connection state from the binary blob, and infuse the connection state into the second protocol stack at the second device..

4. (Canceled)

5. (Currently Amended): The one One or more processor-accessible storage media as recited in claim 4, comprising processor-executable instructions that, when executed, direct a first device to perform actions comprising:

accepting a connection from a connecting device at a forwarder;
receiving data at the forwarder from the connecting device as a result of
accepting the connection;
forwarding the data from the forwarder to a classifier;
determining, by the classifier, a second device for receiving the connection;
aggregating a connection state for the connection at the classifier by
aggregating a protocol state of a first protocol stack and the data to constitute a binary
blob, wherein the aggregating the protocol state comprises compiling the protocol state

from the first protocol stack for use in offloading the connection state as the binary blob, wherein the compiled protocol state includes destination and source ports and IP addresses, wherein the action of compiling further comprises an action of: compiling the protocol state from the first protocol stack starting at a highest level of the first protocol stack, proceeding down the first protocol stack, to compile the protocol state, and then aggregating the received data with the compiled protocol state into the binary blob to be sent to the second device;

sending the connection state from the classifier to the second device for injection into a second protocol stack at the second device by sending the binary blob including the protocol state and the data to the second device, whereby the connection is transferred to the second device;

in conjunction with sending the connection state, adding an entry to a mapping table maintained by the forwarder that indicates the second device as a destination for packets for the connection;

sending a mapping for a flow identifier to the second device based upon the entry in the mapping table;

receiving subsequent communications from the connecting device by the forwarder; and

encapsulating the subsequent communications by the forwarder according to the entry in the mapping table of the forwarder by inserting the flow identifier into the encapsulated communications, wherein the flow identifier serves to identify a flow of encapsulated communications received at the second device from the forwarder as being associated with the connection to the connecting device.

6. (Currently Amended): The one or more processor-accessible storage media as

recited in claim ~~4, 5,~~ wherein the action of compiling comprises an action of:

compiling the protocol state from the first protocol stack at a transmission control protocol (TCP) stack portion and an internet protocol (IP) stack portion.

7. (Previously Presented): The one or more processor-accessible storage media as

recited in claim 1, wherein the action of sending comprises actions of:

bundling the connection state with the mapping for the flow identifier that corresponds to the connection to produce the binary blob; and

transmitting the binary blob having the flow identifier mapping bundled therein from the classifier to the second device.

8. (Canceled)

9. (Previously Presented): The one or more processor-accessible storage media as

recited in claim 1, comprising the processor-executable instructions that, when

executed, direct the classifier to perform further actions comprising:

selecting the flow identifier for the connection responsive to a connection counter; and

sending mapping for the flow identifier to the second device for use by the second device in identifying a source of encapsulated communications received by the second device from the forwarder and corresponding to the connection.

10. (Previously Presented): The one or more processor-accessible storage media as recited in claim 1, wherein the processor-executable instructions, when executed, direct the forwarder to perform a further action comprising:

forwarding subsequent communications for the connection to the second device using the flow identifier to encapsulate the subsequent communications, said encapsulated subsequent communications including the flow identifier in source and destination port fields of a TCP (Transmission Control Protocol) header.

11. (Currently Amended): One or more processor-accessible storage media comprising processor-executable instructions that, when executed, direct a first device and a second device to perform actions comprising:

accepting a connection from a connecting device by a forwarder at the first device;

receiving data at the first device as a result of accepting the connection;

aggregating, by a classifier at the first device, a connection state for the connection at the first device by aggregating a protocol state of a first protocol stack and the received data to constitute an aggregated connection state;

sending the aggregated connection state including the protocol state and the received data asynchronously from the first device to the second device;

receiving the aggregated connection state asynchronously at the second device, whereby the aggregated connection state comprised of the protocol state and the received data is received intact at the second device, wherein the receiving further

comprises receiving the connection state as a binary blob asynchronously at the second device, and recognizing the binary blob as a blob for connection migration;

injecting the aggregated connection state for the connection into a network stack at the second device by infusing the protocol state into a second protocol stack forming a portion of the network stack at the second device, and directing data from the binary blob to an application at the second device as if the second device were part of a new locally terminated connection;

in conjunction with sending the aggregated connection state, sending a mapping for a flow identifier from the first device to the second device, the flow identifier for identifying encapsulated packets received from the forwarder;

continuing the connection at the second device using the injected connection state;

receiving subsequent communications from the connecting device by the forwarder;

encapsulating the subsequent communications by the forwarder by inserting the flow identifier into the encapsulated communications according to a mapping table maintained by the forwarder; and

receiving the encapsulated communications at the second device from the forwarder, wherein the flow identifier serves to identify a flow of encapsulated communications as being associated with the connection to the connecting device according to the mapping for the flow identifier received from the first device.

12. (Previously Presented): The one or more processor-accessible storage media as recited in claim 11, wherein the action of continuing comprises an action of:

continuing the connection by forwarding received packets received by the forwarder to the second device, wherein the second device includes a migrator intermediate driver for buffering packets received prior to the injected aggregated connection state becoming active on the second device.

13. (Canceled)

14. (Previously Presented): The one or more processor-accessible storage media as recited in claim 11, wherein the action of injecting the connection state further comprises an action of:

indicating the data for the connection up the network stack toward an application.

15. (Previously Presented): The one or more processor-accessible storage media as recited in claim 11, wherein the action of injecting comprises an action of:

infusing the protocol state from the connection state into a second protocol stack forming a portion of the network stack.

16. (Previously Presented): The one or more processor-accessible storage media as recited in claim 15, wherein the action of infusing comprises an action of:

infusing the protocol state into the second protocol stack starting at a highest level of the second protocol stack.

17. (Previously Presented): The one or more processor-accessible storage media as recited in claim 11, wherein the action of receiving comprises an action of:

receiving a binary blob from the first device at the second device, the binary blob including the aggregated connection state bundled with the mapping for the flow identifier that corresponds to the connection.

18. (Previously Presented): The one or more processor-accessible storage media as recited in claim 17, wherein the action of receiving comprises actions of:

unbundling the aggregated connection state and the mapping for the flow identifier at a level of the network stack that is below a second protocol stack portion of the network stack.

19. (Previously Presented): The one or more processor-accessible storage media as recited in claim 11, comprising the processor-executable instructions that, when executed, direct the second device to perform further actions comprising:

receiving the mapping for a flow identifier at the second device from the first device; and

storing the received mapping in an encapsulation mapping table that is accessed according to the flow identifier; and

receiving the encapsulated communications at the second device from the forwarder at the first device, said encapsulated communications including the flow identifier in source and destination port fields of a TCP (Transmission Control Protocol) header.

20. (Previously Presented): The one or more processor-accessible storage media as recited in claim 11, wherein the processor-executable instructions, when executed, direct the second device to perform a further action comprising:

receiving from the first device the encapsulated communications that have the flow identifier in source and destination port fields of a TCP (Transmission Control Protocol) header; and

de-encapsulating the encapsulated packets using an encapsulation mapping entry corresponding to the received mapping for the flow identifier that links the flow identifier to a source/destination pair.

21. - 87. (Canceled)

88. (Currently Amended): A method of carrying out load balancing, comprising:

accepting a connection from a connecting device at a forwarder, the forwarder including one or more processors implementing instructions contained in one or more processor-accessible storage media for [::] receiving data at the forwarder from the connecting device as a result of accepting the connection;

forwarding the data from the forwarder to a classifier;

determining, by the classifier, a second device for receiving the connection; aggregating a connection state for the connection at the classifier by aggregating a protocol state of a first protocol stack and the data to constitute a binary blob;

sending the connection state from the classifier to the second device for injection into a second protocol stack at the second device by sending the binary blob including the protocol state and the data to the second device, whereby the connection is transferred to the second device, wherein the sending the connection state further comprises transmitting the binary blob from the classifier to the second device asynchronously via the forwarder in a reliable manner such that the binary blob is received intact at the second device even if one or more packets that comprise the binary blob are lost or corrupted;

in conjunction with sending the connection state, adding an entry to a mapping table maintained by the forwarder that indicates the second device as a destination for packets for the connection;

sending a mapping for a flow identifier to the second device based upon the entry in the mapping table;

receiving subsequent communications from the connecting device by the forwarder; and

encapsulating the subsequent communications by the forwarder according to the entry in the mapping table of the forwarder by inserting the flow identifier into the encapsulated communications; and receiving the encapsulated communications at the second device from the forwarder, wherein the flow identifier serves to identify a flow of

encapsulated communications received at the second device from the forwarder as being associated with the connection to the connecting device.

89. (Currently Amended): A method of carrying out load balancing, comprising:

accepting a connection from a connecting device by a forwarder at a first device, the first device including one or more first processors implementing instructions contained in one or more first processor-accessible storage media for [();]receiving data at the first device as a result of accepting the connection;

aggregating, by a classifier at the first device, a connection state for the connection at the first device by aggregating a protocol state of a first protocol stack and the received data to constitute an aggregated connection state;

sending the aggregated connection state including the protocol state and the received data asynchronously from the first device to a second device, the second device including one or more second processors implementing instructions contained in one or more second processor-accessible storage media for [();]receiving the aggregated connection state asynchronously at the second device, whereby the aggregated connection state comprised of the protocol state and the received data is received intact at the second device, wherein the receiving further comprises receiving the aggregated connection state as a binary blob asynchronously at the second device, and recognizing the binary blob as a blob for connection migration;

injecting the aggregated connection state for the connection into a network stack at the second device by infusing the protocol state into a second protocol stack forming a portion of the network stack at the second device, and directing data from the

binary blob to an application at the second device as if the second device were part of a new locally terminated connection;

in conjunction with sending the aggregated connection state, sending a mapping for a flow identifier from the first device to the second device, the flow identifier for identifying encapsulated packets received from the forwarder;

continuing the connection at the second device using the injected connection state;

receiving subsequent communications from the connecting device by the forwarder;

encapsulating the subsequent communications by the forwarder by inserting the flow identifier into the encapsulated communications according to a mapping table maintained by the forwarder; and

receiving the encapsulated communications at the second device from the forwarder, wherein the flow identifier serves to identify a flow of encapsulated communications as being associated with the connection to the connecting device according to the mapping for the flow identifier received from the first device.

90. (Canceled)

91. (New): A method of carrying out load balancing, the method comprising:

accepting a connection from a connecting device at a forwarder, the forwarder including one or more processors implementing instructions contained in one or more

processor-accessible storage media for receiving data at the forwarder from the connecting device as a result of accepting the connection;

forwarding the data from the forwarder to a classifier;

determining, by the classifier, a second device for receiving the connection;

aggregating a connection state for the connection at the classifier by aggregating a protocol state of a first protocol stack and the data to constitute a binary blob, wherein the aggregating the protocol state comprises compiling the protocol state from the first protocol stack for use in offloading the connection state as the binary blob, wherein the compiled protocol state includes destination and source ports and IP addresses, wherein the compiling further comprises compiling the protocol state from the first protocol stack starting at a highest level of the first protocol stack, proceeding down the first protocol stack, to compile the protocol state, and then aggregating the received data with the compiled protocol state into the binary blob to be sent to the second device;

sending the connection state from the classifier to the second device for injection into a second protocol stack at the second device by sending the binary blob including the protocol state and the data to the second device, whereby the connection is transferred to the second device;

in conjunction with sending the connection state, adding an entry to a mapping table maintained by the forwarder that indicates the second device as a destination for packets for the connection;

sending a mapping for a flow identifier to the second device based upon the entry in the mapping table;

receiving subsequent communications from the connecting device by the forwarder; and

encapsulating the subsequent communications by the forwarder according to the entry in the mapping table of the forwarder by inserting the flow identifier into the encapsulated communications, wherein the flow identifier serves to identify a flow of encapsulated communications received at the second device from the forwarder as being associated with the connection to the connecting device.